



University of Dundee

Computing canonical images in permutation groups with Graph Backtracking

Jefferson, Christopher; Waldecker, Rebecca; Wilson, Wilf A.

DOI:
[10.1016/j.jaca.2023.100010](https://doi.org/10.1016/j.jaca.2023.100010)

Publication date:
2023

Licence:
CC BY

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Discovery Research Portal](#)


Citation for published version (APA):
Jefferson, C., Waldecker, R., & Wilson, W. A. (2023). Computing canonical images in permutation groups with Graph Backtracking. *Journal of Computational Algebra*, 8, Article 100010.
<https://doi.org/10.1016/j.jaca.2023.100010>

General rights

Copyright and moral rights for the publications made accessible in Discovery Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Computing canonical images in permutation groups with Graph Backtracking

Christopher Jefferson, Rebecca Waldecker*, Wilf A. Wilson

ARTICLE INFO

Keywords:

Permutation groups
Search algorithms
Canonical image
Backtrack methods

ABSTRACT

We describe a new algorithm for finding a canonical image of an object under the action of a finite permutation group. This algorithm builds on previous work using Graph Backtracking [9], which extends Jeffrey Leon's Partition Backtrack framework [14,15]. Our methods generalise both Nauty [17] and Steve Linton's Minimal image algorithm [16].

1. Introduction and background

Given a group G that acts on a finite set Ω , one of the simplest questions we can ask is whether two elements $\omega_1, \omega_2 \in \Omega$ are in the same orbit of G , and if so, to find an element $g \in G$ such that $\omega_1^g = \omega_2$.

This problem turns out to be surprisingly difficult. If we consider Ω as the set of graphs on a vertex set V , and $G = \text{Sym}(V)$, then the above question is equivalent to the graph isomorphism problem, which has been studied extensively both in theory [1] and in practice. The best-known practical system for solving the graph isomorphism problem is Nauty [17]. The way in which Nauty solves this problem in practice, which we will also refer to in this article, is to use a **canonical labelling function**. Given a group G acting on a set Ω , a canonical labelling function maps each element $\omega \in \Omega$ to a "canonical" member $\text{CANONICAL}(\omega)$ of its G -orbit, which we call its canonical image. This means that two elements of Ω are in the same G -orbit if and only if they have the same canonical image, i.e. the same image under the canonical labelling function.

Canonical labelling functions have several applications. As a brief example, suppose that we are given a set of graphs, all with vertex set Ω , and we want to decide for any two given graphs whether or not they are in the same $\text{Sym}(\Omega)$ -orbit with respect to the naturally induced action. Given only two graphs, we could use a function which directly checks if the two graphs are in the same orbit, which may be faster than finding canonical images. However, for larger sets of graphs, if we can only check whether or not two graphs are in the same orbit, and if all the graphs are in different orbits, then we need to compare every graph to every other graph in order to prove that no two different graphs are in same orbit. If, instead, we can compute the canonical

image of a graph under the action of $\text{Sym}(\Omega)$, then we only need to look at each graph once, find its canonical image, and then we know that the graphs in the same orbit are exactly the ones with the same canonical image.

In previous work, we have discussed methods for calculating minimal and canonical images (see [8]), building on the work of Linton ([16]). Minimal images are a special case of canonical images, using a total ordering of the underlying set. This previous work was limited to the induced action of a permutation group on subsets of a set. Another important special case is finding canonical images of graphs (Nauty, see [17]). For completeness, we mention the related but very different problem of finding a canonical member of every orbit under a given group action. This problem is treated in [2]. There is also related work by Schweitzer and Wiebking on canonical images from a more general perspective (see [18]), but we do not discuss it here. As far as we understand, their work considers the problem from a theoretical perspective and is not intended to lead to a practical algorithm, and in particular we cannot compare our methods to theirs in terms of performance.

In this paper we describe a general and practical framework for finding canonical images. Our results build on Graph Backtracking ([9] and [12]), which itself builds on Nauty and on the search method developed by Jeffrey Leon, known as Partition Backtracking. The use of the Graph Backtracking framework implies that we will re-use some concepts that we have introduced in previous articles, but with minor changes and with a stronger focus on practical implementation. While it is helpful to be familiar with [9] when reading this article, it is not absolutely necessary, because we state the most important definitions and results and explain how we adapted them for this new work.

* Corresponding author.

E-mail address: rebecca.waldecker@mathematik.uni-halle.de (R. Waldecker).

We begin with general comments and some examples for motivation, set-up and notation in Section 2. There we also adjust some earlier definitions for the purpose of this article, and introduce some new definitions. In Section 3 we explain the core of our technique, and how we build search trees, again with examples for illustration, and then in Section 4 we explain how we construct canonical images from such search trees. Section 5 discusses some experiments, including comparisons with earlier work. We close with a few final remarks and an appendix with some additional comments and results, including further details relevant for implementation.

2. Preliminaries and first examples

Throughout this paper Ω is a finite set and G is a subgroup of $\text{Sym}(\Omega)$. In our applications we often consider G in its induced action on sets of lists, graphs with vertex set Ω , subsets of Ω , subgroups of G (like point stabilisers) or other combinatorial structures. We use the name ACTSET for a set of combinatorial objects on which G acts and we will find canonical images for members of ACTSET in their orbit in G . Throughout, we use standard notation for permutation groups and group actions without further explanation, in particular notation for orbits and for the stabilisers of points, lists or sets. See for example [4].

In order to efficiently determine whether two elements of ACTSET are in the same orbit of G , we apply the concept of a canonical image (see also Definition 2.1 from [8]).

Definition 2.1. A **canonical labelling function** C with respect to the action of G on a set ACTSET is a function $C : \text{ACTSET} \mapsto \text{ACTSET}$ such that, for all $a \in \text{ACTSET}$, it is true that:

- $C(a) \in a^G$, and
- $C(a^g) = C(a)$ for all $g \in G$

We call $C(a)$ the **canonical image** of a (with respect to G and the action of G on ACTSET). Whenever the group G and its action on ACTSET are clear from context, then we will not mention them. Further, we say that $g \in G$ is a **canonising element** for a if and only if $a^g = C(a)$.

Inspired by work of Linton [16], we introduced practical algorithms for minimal and canonical images in [8]. In the present paper, we combine these ideas with the concept of Graph Backtrack search, and we give more explanations of implementation details.

One example of a canonical labelling function that is quite intuitive and easy to explain is given by the **minimal image**. It is calculated with respect to a given total ordering \preceq on ACTSET , and it is defined as follows:

$$\text{For all } a \in \text{ACTSET}, \text{Min}_{\preceq}(a) := \min_{\preceq} \{a^g \mid g \in G\}.$$

Unfortunately, for a general canonical image, there is no concise and simple mathematical description of how they are created, because the resulting canonical images are affected by many minor parts of the algorithm. But we can describe the idea in a nutshell: We narrow down the possibilities for a canonical image, efficiently and respecting the action of G , and only at the very end we introduce an ordering and calculate a minimal image. It is only in the last step where details like the exact orderings that are used on Ω and ACTSET influence the outcome of the algorithm. While even this might sound mildly concerning, or at least unsatisfying, this is what happens in practice all along, even for algorithms like Nauty. There, the exact implementation of equitable partitioning determines the canonical graph that is produced. In our algorithms, if the orderings chosen at the end are exactly the same, then the same canonical image will be calculated. In contrast, a canonising element may not be unique. More precisely, given a canonising element $g \in G$ for some object $a \in \text{ACTSET}$, it follows that every element in the right coset $G_a \cdot g$ is a canonising element for a as well.

We close the introduction to this section with two examples that have motivated our work.

Example 2.2.

(a) Nauty finds canonical images of undirected graphs. In our notation, with our given set Ω , this means that $G = \text{Sym}(\Omega)$ and we take ACTSET to be the set of all undirected graphs with vertex set Ω . Then there is a naturally induced action of G on ACTSET with well-defined orbits, and therefore we can apply our notion of a canonical labelling function.

(b) The canonical image algorithm in [8], which extends earlier work of Linton [16], considers any subgroup $G \leq \text{Sym}(\Omega)$, and ACTSET is the set of all subsets of Ω . If $A \subseteq \Omega$, then $g \in G$ maps A to $A^g := \{a^g \mid a \in A\}$. With this induced action, we have well-defined G -orbits again and we can ask for canonical images of elements of ACTSET .

While the two special cases from the previous example are important and occur frequently in practice, there are many other problems we might want to solve. As we will see later, many problems can be transformed and can then be solved as a classical “graph problem” or a classical “canonical image problem”. But there are limits to this approach and it does not always lead to practically useful algorithms, as can be seen in [8].

2.1. Definitions and notation

We keep this section as short as possible – on the one hand we do not want to reproduce all the basic results and examples from previous papers, and on the other hand we do not expect all readers to have read [8] and [9], which is why we give a basic overview:

A **graph** with vertex set Ω is a pair (Ω, E) , where E is a set of 2-subsets of Ω . A **directed graph** with vertex set Ω , or **digraph** for short, is a pair (Ω, A) , where $A \subseteq \Omega \times \Omega$ is a set of pairs of elements in Ω called **arcs**. The elements of Ω are called **vertices** in the context of graphs and digraphs. Our definition allows a digraph to have **loops**, which are arcs of the form (α, α) for some vertex $\alpha \in \Omega$.

Our techniques for searching in $\text{Sym}(\Omega)$ are built around digraphs in which each vertex and arc are given a *label* from a label set \mathfrak{L} . We define a **vertex- and arc-labelled digraph**, or **labelled digraph** for short, to be a triple $(\Omega, A, \text{LABEL})$, where (Ω, A) is a digraph and LABEL is a function from $\Omega \cup A$ to \mathfrak{L} . We often refer to labels for vertices as colours and say that vertices are coloured or individualized.

Let $\text{LABELLED DIGRAPHS}(\Omega, \mathfrak{L})$ denote the class of all labelled digraphs on Ω with labels in \mathfrak{L} . A **labelled digraph stack** on Ω is a finite (possibly empty) list of labelled digraphs on Ω . We denote the collection of all labelled digraph stacks on Ω by $\text{DIGRAPHSTACKS}(\Omega)$. The **length** of a labelled digraph stack S , written $|S|$, is the number of entries that it contains. A labelled digraph stack of length 0 is called **empty**, and we denote the empty labelled digraph stack on Ω by $\text{EMPTYSTACK}(\Omega)$. We use a notation typical for lists, whereby if $i \in \{1, \dots, |S|\}$, then $S[i]$ denotes the i^{th} labelled digraph in the stack S . If $S, T \in \text{DIGRAPHSTACKS}(\Omega)$, then we write $S \cong T$ if and only if S and T are **isomorphic**, which means that they have the same length and there exists a permutation $g \in \text{Sym}(\Omega)$ such that for all $i \in \{1, \dots, |S|\}$, $S[i]^g = T[i]$. The set of isomorphisms from S to T is denoted by $\text{Iso}(S, T)$, and along the same lines, the automorphism group of the stack S is denoted by $\text{Aut}(S)$.

We allow any labelled digraph stack on Ω to be appended onto the end of another. If $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ have lengths k and l , respectively, then we define $S||T$ to be the labelled digraph stack $[S[1], \dots, S[k], T[1], \dots, T[l]]$ of length $k+l$ formed by appending T to S .

Now we discuss some concepts that have appeared in earlier work, but that we will refine and specify for the algorithms relevant to this article. The main reason we have to adjust the existing definitions is that the proofs in [9] and [10] are only concerned with the behaviour of a single search – in standard Graph Backtracking a single problem is solved, and a group, coset, or the empty set is found. When finding canonical images, we must prove that searches for the canonical image of two elements of ACTSET in the same orbit of G produce the same

canonical image and therefore, in several places, we need to add extra conditions about the compatibility with the action of G or $\text{Sym}(\Omega)$.

The following definitions are in the spirit of [9], but they are often adapted for the purpose of this article. There will be several examples in Section 2.2 and later.

Definition 2.3. If $a \in \text{ACTSET}$, then a **refiner for** a is defined to be a function REF from $\text{DIGRAPHSTACKS}(\Omega)$ to itself such that, whenever $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ and $S \cong T$, it is true that

$$\text{Sym}(\Omega)_a \cap \text{Iso}(S, T) \subseteq \text{Sym}(\Omega)_a \cap \text{Iso}(\text{REF}(S), \text{REF}(T)).$$

We say that a refiner REF for an object $a \in \text{ACTSET}$ is **compatible with** $\text{Sym}(\Omega)$ if and only if for all $S \in \text{DIGRAPHSTACKS}(\Omega)$ and all $g \in \text{Sym}(\Omega)$, we have that $\text{REF}(S^g) = (\text{REF}(S))^g$.

The main difference between Definition 2.3 and the refiners from [9] is that we now define refiners for elements a of ACTSET , and then use the group $\text{Sym}(\Omega)_a$, rather than restricting to subgroups of $\text{Sym}(\Omega)$ directly. The additional compatibility condition is necessary here – it captures the fact that refiners for two objects that are in the same orbit must have refiners which operate in “the same way” – this will ensure we find the same canonical image for the two objects. Compatibility will therefore be used in the proofs of the correctness of our algorithms in Sections 3 and 4.

Often, compatibility is automatically fulfilled, as is the case for most of our previously published refiners. But since we did not introduce this notion in earlier work (because it was not needed), we did not check our refiners in [9] for compatibility and therefore we discuss it here. It will turn out that several related concepts must be slightly adapted as well.

When performing our backtracking search, we need to find a way of splitting the current search node into multiple child nodes. In Graph Backtracking, the function used to split has relatively few restrictions, apart from ensuring that the entire space of possible solutions is partitioned between the new child nodes and each child represents a strictly smaller part of the search than its parent.

To ensure that we find the same canonical image when looking at different objects with the same canonical image, we need to more tightly constrain how we split search nodes, taking the action of the group $\text{Sym}(\Omega)$ on ACTSET into account. Once more we adapt definitions from [9]:

Definition 2.4. A **compatible isomorphism approximator** is a function APPROX which maps a pair of labelled digraph stacks on Ω to a subgroup of $\text{Sym}(\Omega)$, a coset of a subgroup of $\text{Sym}(\Omega)$, or the empty set, in such a way that the following conditions are satisfied for all $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ and $g \in \text{Sym}(\Omega)$:

- (i) $\text{Iso}(S, T) \subseteq \text{APPROX}(S, T)$.
- (ii) If $|S| \neq |T|$, then $\text{APPROX}(S, T) = \emptyset$.
- (iii) If $\text{APPROX}(S, T) \neq \emptyset$, then there is some $h \in \text{Sym}(\Omega)$ such that $\text{APPROX}(S, T) = \text{APPROX}(S, S) \cdot h$.
- (iv) $\text{APPROX}(S^g, T^g) = \text{APPROX}(S, T)^g$.

Given $S \in \text{DIGRAPHSTACKS}(\Omega)$, we usually abbreviate $\text{APPROX}(S, S)$ as $\text{APPROX}(S)$.

The isomorphism approximators defined in Section 5 of [9] are all compatible already. The first, given in Section 5.1, calculates $\text{Iso}(S, T)$ exactly, hence it is naturally compatible: For all graph stacks S and T and all permutations $g \in \text{Sym}(\Omega)$, it is true that $\text{Iso}(S^g, T^g) = \text{Iso}(S, T)^g$.

Another way of viewing Condition (iv) of Definition 2.4 is that the isomorphism approximator commutes with re-naming the elements of Ω . Using this observation, we can see that the other approximators from Section 5 of [9] are also compatible, because none of the approximations ever refer to an ordering, or particular elements, of Ω . We next

similarly extend the definition of a splitter from Section 6 of [9], again adding an extra requirement to ensure compatibility.

Definition 2.5. Suppose that APPROX is a compatible isomorphism approximator and let \mathcal{T} denote the set of labelled digraph stacks T on Ω where $|\text{APPROX}(T)| > 1$. Let SPLIT be a function from \mathcal{T} to the set of finite ordered lists of labelled digraph stacks on Ω .

We say that SPLIT is a **compatible splitter** (with respect to APPROX) if and only if, for all $T \in \mathcal{T}$ and $n \in \mathbb{N}$, the ordered list $\text{SPLIT}(T) = [T_1, \dots, T_n]$ satisfies the following:

- (i) $\text{Iso}(T, T) = \text{Iso}(T \parallel T_1, T \parallel T_1) \cup \text{Iso}(T \parallel T_1, T \parallel T_2) \cup \dots \cup \text{Iso}(T \parallel T_1, T \parallel T_n)$.
- (ii) $|\text{APPROX}(T \parallel T_i)| < |\text{APPROX}(T)|$ for all $i \in \{1, \dots, n\}$.
- (iii) For all $g \in \text{Sym}(\Omega)$, the sets $\{T_1, \dots, T_n\}^g$ and $\{T_1^g, \dots, T_n^g\}$ coincide. This means that the lists $\text{SPLIT}(T^g)$ and $(\text{SPLIT}(T))^g$ are equal up to ordering.

How can we create compatible splitters? The “fixed point splitter” of [9, Definition 6.4], does not satisfy our requirements, because it takes the smallest member of Ω which is in a shortest orbit (ignoring fixed points). This will usually not be compatible with the application of permutations. In the next subsection, we build some necessary machinery to create compatible splitters.

2.2. Approximators for canonisation

To create compatible splitters, we first introduce a way of ordering (our current approximation of) the orbits of the group.

Definition 2.6. An **ordered orbit estimator** for labelled digraph stacks is a function PART that maps a labelled digraph stack on Ω to an ordered partition of Ω such that for each $S \in \text{DIGRAPHSTACKS}(\Omega)$ the following holds:

- If $i, j \in \Omega$ are in the same orbit of $\text{Aut}(S)$, then they are in the same cell of $\text{PART}(S)$, and
- $\text{PART}(S)^g = \text{PART}(S^g)$ for all $g \in \text{Sym}(\Omega)$.

From there we can define an **ordered orbit approximator** APPROX_O as follows: Given two labelled digraph stacks $S, T \in \text{DIGRAPHSTACKS}(\Omega)$, we define $\text{APPROX}_O(S, T)$ as the set of permutations in $\text{Sym}(\Omega)$ that map $\text{PART}(S)$ to $\text{PART}(T)$. As before, if $S \in \text{DIGRAPHSTACKS}(\Omega)$, then we abbreviate $\text{APPROX}_O(S, S)$ as $\text{APPROX}_O(S)$.

The first property in our definition of an ordered orbit estimator can be expressed as “ $\text{PART}(S)$ is (possibly) coarser than the orbit partition of $\text{Aut}(S)$ ”. The reason why we allow a coarser partition is that calculating the exact orbits of $\text{Aut}(S)$ would require solving a graph isomorphism problem. For the corresponding approximator, we note that just as for isomorphism approximators, the images under this map are subgroups of $\text{Sym}(\Omega)$, cosets of a subgroup or the empty set. Next we give examples and we show how an ordered orbit estimator may merge together different $\text{Aut}(S)$ -orbits:

Example 2.7. Suppose that $\Omega := \{1, 2, 3, 4, 5, 6, 7, 8\}$ and that Γ_1 is the graph with vertex set Ω and the following edges (viewed as arcs both ways): $\{1, 2\}, \{1, 3\}, \{2, 3\}, \{5, 6\}, \{6, 7\}$. Let S be the digraph stack with just the entry Γ_1 .

Also, let Γ_2 be the graph with vertex set Ω and edges $\{1, 6\}, \{1, 7\}, \{2, 6\}, \{2, 7\}, \{3, 8\}, \{4, 8\}$, and $\{5, 8\}$, again viewed as arcs both ways. Let T be the digraph stack with just the entry Γ_2 .

(a) Let F denote the function that, for each graph in the stack, creates an ordered partition where the cells correspond to the number of neighbours that each vertex has. In practice, we need to decide in which

order to organise the cells and apply this ordering consistently. For example, S could get mapped to $F(S) := [4, 8|5, 7|1, 2, 3, 6]$ and T could get mapped to $F(T) := [3, 4, 5|1, 2, 6, 7|8]$. Just to clarify our notation for ordered partitions: The vertical lines separate the cells, and the cells themselves are considered to be sets (hence unordered), while the ordering of the cells within the partition matters.

In order to check, at least for these two stacks, whether or not F is an ordered orbit estimator in the sense of Definition 2.6, we need to consider the orbits of $\text{Aut}(S)$, namely $\{1, 2, 3\}$, $\{4, 8\}$, $\{5, 7\}$, and $\{6\}$, and the orbits of $\text{Aut}(T)$, which are $\{1, 2, 6, 7\}$, $\{3, 4, 5\}$, and $\{8\}$.

A simple check shows that all elements that are in the same orbit under the automorphism group also are in the same cell. Also, the number of neighbours of vertices is invariant under the action of $\text{Sym}(\Omega)$, so that both conditions for an ordered orbit estimator are satisfied. The image $\text{APPROX}_O(S, T)$ under the corresponding approximator is the empty set, because there is no element in $\text{Sym}(\Omega)$ that maps $F(S)$ to $F(T)$.

(b) Another possibility would be the function F that, for each graph in the stack, creates an ordered partition where the cells correspond to connected components. Again, in practice, we need to decide how to order the cells. For example $F(S) = [4|8|1, 2, 3|5, 6, 7]$ and $F(T) = [1, 2, 6, 7|3, 4, 5, 8]$. We do our simple check, and this time we do not obtain an ordered orbit estimator. The isolated vertices 4 and 8 in the graph Γ_1 are in the same $\text{Aut}(S)$ -orbit, but in different cells. For Γ_2 , we see that the vertex 8 is not in the same $\text{Aut}(T)$ -orbit as 3, 4 and 5, but it is in the same cell of $F(T)$. Therefore, taking the connected components does not create a valid ordered orbit estimator. However, taking connected components commutes with the action of $\text{Sym}(\Omega)$, it is therefore not an unnatural way to produce a partition.

Rephrasing the first condition of an ordered orbit estimator gives another way to look at this: If $S \in \text{DIGRAPHSTACKS}(\Omega)$, then $\text{Aut}(S)$ is a (possibly proper) subgroup of the stabiliser of the ordered partition $\text{PART}(S)$ in $\text{Sym}(\Omega)$. In that light, connected components produce too fine a partition, because members of the same orbit of the automorphism can be in different connected components. While the automorphism group can permute connected components, it cannot permute its own orbits non-trivially.

Using ordered orbit estimators and approximators, we can create a compatible splitter, as demonstrated in the next example:

Example 2.8 (Compatible splitter).

Suppose that E is an ordered orbit estimator on $\text{DIGRAPHSTACKS}(\Omega)$ and that APPROX_E denotes the corresponding approximator as in Definition 2.7. Now we define a compatible splitter SPLIT_E as follows: For any $S \in \text{DIGRAPHSTACKS}(\Omega)$, calculate $E(S)$. If all cells are singletons, then we do not consider S because there is a unique permutation in $\text{Sym}(\Omega)$ that maps $E(S)$ to itself, and hence $|E(S)| = 1$.

Otherwise we determine the size k_S of the shortest non-singleton cell and we take the first cell in $E(S)$ of this minimal size k_S , and then generate a list $[S_1, \dots, S_{k_S}]$ of graph stacks, each consisting of just one graph. For every element in this shortest non-singleton cell, this unique graph in the stack is such that the vertex for this element is individualised. This means the graph has no edges, all vertices except the individualised vertex are labelled 0, and the individualised vertex is labelled 1. (The exact labels are unimportant, but the same labels must always be used.) The automorphism group of this graph fixes the individualised vertex and allows all other vertices to freely permute. Finally, we set $\text{SPLIT}_E(S) := [S_1, \dots, S_{k_S}]$ as described above.

Going back to Example 2.7, we make this explicit and look at the ordered orbit estimator F that we introduced in (a). The images were $F(S) = [4, 8|5, 7|1, 2, 3, 6]$ and $F(T) = [3, 4, 5|1, 2, 6, 7|8]$.

For each number $i \in \{1, \dots, 8\}$, let Δ_i denote the stack that consists solely of the empty graph on $\{1, \dots, 8\}$ where the vertex i is individualised as described above. Then the ordering of the cells that was used for F gives that $\text{SPLIT}_F(S) = [\Delta_4, \Delta_8]$, because 2 is the size of

the shortest non-singleton cell and the first such cell is $\{4, 8\}$, and $\text{SPLIT}_F(T) = [\Delta_3, \Delta_4, \Delta_5]$, because here we have a unique shortest non-singleton cell, namely $\{3, 4, 5\}$.

In some situations we will need a fixed point approximator, which was discussed in our previous work and which gives a subset of the fixed points set of the graph stack. This can be easily created from an ordered partition approximator:

Definition 2.9. We define the function FIX from the set of all ordered partitions on Ω to the set of all ordered lists with entries from Ω as follows:

For each ordered partition P of Ω , $\text{FIX}(P)$ is defined to be the list of elements of P that are in singleton cells, in the order in which they appear in P .

As an example for this last definition, we look at $\Omega := \{1, \dots, 8\}$ and we consider $P := [4|8|1, 2, 3|6|5, 7]$. Then $\text{FIX}(P) = [4, 8, 6]$.

In line with the compatibility conditions that we have introduced so far, we introduce a compatible partial order, which we will use for ordering the nodes of our search trees.

Definition 2.10. A compatible partial order on the set of digraph stacks on Ω is a partial order \leq on $\text{DIGRAPHSTACKS}(\Omega)$ such that, for all $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ and all $g_1, g_2 \in \text{Sym}(\Omega)$:

$$S \leq T \text{ if and only if } S^{g_1} \leq T^{g_2}.$$

For the next definition we recall that the images of an ordered orbit estimator are ordered partitions of Ω , and we give an example of a compatible partial order. It will be discussed in Lemma 2.13.

Definition 2.11. Suppose that PART is an ordered orbit estimator on $\text{DIGRAPHSTACKS}(\Omega)$ and let $S, T \in \text{DIGRAPHSTACKS}(\Omega)$.

We define $S \leq T$ if and only if one of the following holds: The list $[|O| \mid O \in \text{PART}(S)]$ is shorter than the list $[|O| \mid O \in \text{PART}(T)]$, or the lists have the same length and $[|O| \mid O \in \text{PART}(S)]$ comes first under lexicographic ordering.

Example 2.12. We build on Example 2.7, and again we let $\Omega := \{1, 2, 3, 4, 5, 6, 7, 8\}$ and we use the graphs Γ_1 and Γ_2 as well as the stacks S and T from that example, along with the function F that, for each graph in the stack, creates an ordered partition where the cells correspond to the number of neighbours that each vertex has. We recall that $F(S) = [4|8|5, 7|1, 2, 3, 6]$ and $F(T) = [3, 4, 5|1, 2, 6, 7|8]$ once we decided that we begin with the lowest number of neighbours.

In order to compare S and T with \leq from the previous definition, we write down the lists of cell sizes. For $F(S)$, we obtain $[1, 1, 2, 4]$, and for $F(T)$, we obtain $[3, 4, 1]$, which means that $T \leq S$. Now we check for compatibility, hence we let $g_1, g_2 \in \text{Sym}(\Omega)$. Given the action of $\text{Sym}(\Omega)$ on $\text{LABELLED DIGRAPHS}(\Omega, \mathcal{Q})$, it follows that $F(S^{g_1}) = F(S)$ if we order the cells in the same way, starting with the lowest number of neighbours. Also $F(T^{g_2}) = F(T)$, for the same reason, and consequently $T^{g_2} \leq S^{g_1}$. Therefore we have an example of a compatible partial order, and this is in fact the ordering we will use later on.

The following lemma shows that the construction from the previous example works more generally:

Lemma 2.13. Let PART be an ordered orbit estimator on $\text{DIGRAPHSTACKS}(\Omega)$ and let \leq be as defined in Definition 2.11. Then this order is a compatible partial order.

Proof. We have seen the main argument in Example 2.12. Let $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ be such that $S \leq T$ and let $g_1, g_2 \in \text{Sym}(\Omega)$.

The compatibility of PART implies that $\text{PART}(S)^{g_1} = \text{PART}(S^{g_1})$, and it follows that

$$[|O| \mid O \in \text{PART}(S^{g_1})] = [|O| \mid O \in \text{PART}(S)^{g_1}] = [|O| \mid O \in \text{PART}(S)].$$

This is because the action of permutations on ordered partitions respects the ordering and size of the cells. For the same reason we see that $[|O| \mid O \in \text{PART}(T^{g_2})] = [|O| \mid O \in \text{PART}(T)]$, which implies that $S^{g_1} \leq T^{g_2}$. Since permutations are bijective, we can argue in the same way for the reverse implication. \square

3. Searching with graph stacks

Now that we have introduced compatible refiners and compatible splitters, we will show how we use these together to search for candidates for canonical images. This search is split into two parts. The present section will build most of the machinery, using a backtrack search with graph stacks and then return a list of candidates. The next section will show how to systematically pick a canonical image from this list of candidates. While this algorithm uses much of the same machinery as Graph Backtracking, it behaves quite differently. Most notably, in Graph Backtracking there is always a pair of graph stacks, called the left and right hand stacks, and the search stops if there are no isomorphisms from the left hand stack to the right. When looking for canonical images, we will only consider a *single* graph stack. This stack corresponds to the *right hand side stack in Graph Backtracking*.

3.1. Building search trees

Suppose that, for some finite set Ω , some group $G \leq \text{Sym}(\Omega)$ and some set of objects ACTSET, we wish to find canonical images for members of ACTSET with respect to the action of G . Our algorithm requires the following: a set of compatible refiners for the elements of ACTSET that we wish to find canonical images for (only one of these will be used during any invocation of the algorithm), a compatible splitter SPLIT, an ordered orbit estimator PART and corresponding approximator APPROX, a refiner R_G for G . For R_G we use the original definition of a refiner from [9], because R_G does not need to be compatible. There are several different choices for R_G , for example we can use the orbits of G . We will discuss more options in the appendix of this paper.

All the concepts mentioned above will be used throughout this section without reintroduction. Note that changing any of the refiners, the splitter, or the ordered orbit estimator and approximator can change the resulting canonical image.

We begin by giving a slightly simplified explanation of our method in Algorithm 1. Note that this initial presentation shows an algorithm which correctly returns a set of permutations, namely a set of candidates for permutations that map $a \in \text{ACTSET}$ to a canonical image. However, this set is produced in a very inefficient way. We will then show how the search can be improved. This second algorithm uses the function MINPERM which, given a group G and a list of integers L , returns a permutation $g \in G$ which maps L point-wise to its lexicographic minimal image. An implementation of this function is given in Section 7.

We can view an execution of Algorithm 1 as a search tree (Definition 3.1).

Definition 3.1. A finite directed graph is a **search tree** if and only if it has the following properties:

- It is connected,
- it has a unique initial node (the **root node**) with no incoming arcs and only outgoing arcs,
- every node except for the root node has exactly one incoming arc, and
- the graph contains no cycles (neither directed nor undirected).

This means that for every node, there is a unique directed path in the graph from the root node to it. The **children** of a node v are all nodes

Algorithm 1 A recursive algorithm using labelled digraph stacks to search in $\text{Sym}(\Omega)$.

Input: $a \in \text{ACTSET}$ and a compatible refiner R_a for a , taken from a set of compatible refiners for ACTSET.

Output: A set of permutations $g \in G$ containing the permutations which map a to its canonical representative.

```

1: procedure SEARCH( $S$ )                                ▷ The main recursive search procedure
2:    $S' \leftarrow \text{REFINE}(S)$                                ▷ Refine the stack
3:   if  $|\text{APPROX}(S')| = 1$  then
4:     return  $\text{MINPERM}(G, \text{FIX}(\text{PART}(S')))$ 
5:   else                                               ▷ Multiple potential solutions
6:      $[S_1, \dots, S_l] \leftarrow \text{SPLIT}(S')$              ▷ Split the search space
7:     return  $\bigcup_{i \in \{1, \dots, l\}} \text{SEARCH}(S'_i \| S_i)$    ▷ Search recursively

8: procedure REFINE( $S$ )                                ▷ Attempt to prune the search space
9:   while TRUE do                                     ▷ Loop until fixed point
10:     $S_C \leftarrow S$                                    ▷ Save the stack to check if progress is made
11:    for  $R \in \{R_a, R_G\}$  do
12:       $S \leftarrow S \| R(S)$                            ▷ Apply each refiner in turn
13:    if  $|\text{APPROX}(S)| < |\text{APPROX}(S_C)|$  then
14:      return  $S$                                        ▷ Stop: the refinement is no longer making progress

15: return  $\text{SEARCH}(\text{EMPTYSTACK}(\Omega))$                  ▷ Start the search

```

which are connected to an outgoing arc of v . The nodes that only have an incoming arc and no outgoing arcs will be called **leaves**.

We allow a search tree to be vertex-labelled and we place no restrictions on labels used on vertices.

We now describe Algorithm 1 in terms of a search tree by creating a node for each call of SEARCH(S), labelling this node with the *graph stack* S . We define a directed edge from any node to the nodes which are created by the recursive calls to SEARCH Line 7 from that node. Leaf nodes (which originate from calls to SEARCH(S) which return on Line 4) receive an additional label, namely a **result** permutation, which is the permutation returned on Line 4.

In the next section we describe how to find a canonical image from the search tree. We will not reproduce all results and definitions from our previous work, for more technical details see [9] and [11]. We now give a full example of Algorithm 1. This example does not demonstrate all possible cases. For a more complete discussion of applying refiners during search, and the issues involved, see [9].

Example 3.2. In this example, $\Omega := \{1, \dots, 8\}$, $G := \langle (1, 2, 3, 4, 5, 6, 8), (1, 3, 2, 6, 4, 5), (1, 6)(2, 3)(4, 5)(7, 8) \rangle \leq \text{Sym}(\Omega)$ and our object a from ACTSET is a graph Γ with vertex set Ω and list of edges $\{1, 6\}, \{1, 7\}, \{2, 6\}, \{2, 7\}, \{3, 8\}, \{4, 8\}, \{5, 8\}$. For our approximator we use APPROX $_G$ from Definition 5.9 of [9], and we explain explicitly what it does in this specific example. The refiners that we use are R_a and R_G . For R_a , we convert Γ into a digraph by replacing each edge with arcs in both directions, and then we add this graph to the stack (if it is not already in the stack). The refiner R_G adds a graph that connects vertices, again with arcs in both directions, if and only if they are in the same G -orbit.

Now we can build a search tree for Γ as described in Algorithm 1. At the first call of SEARCH (which is also the root node), S is the empty stack. We begin by calling REFINE. We call R_a and R_G , and as explained above this means that we add a directed version of Γ and then a graph whose connected components are exactly the G -orbits.

We then calculate the image under the ordered orbit approximator APPROX, as we have seen in Example 2.7 (a). There are three vertices in the graph a with exactly one neighbour (namely 3, 4 and 5), four vertices with two neighbours (namely 1, 2, 6, 7) and there is a unique vertex with three neighbours (namely 8). For our bookkeeping in this example we write the image down as an ordered partition of Ω : $P_1 = [3, 4, 5]1, 2, 6, 7|8$. Note that the refiner for R_a does not add anything anymore, because the graph stack already contains a directed version of the original graph a , so we will not discuss it further in this example.

We call R_G again, using the fact that 8 is the only singleton cell, so we can consider the orbits of G_8 . This does not produce any new information, because G is 2-transitive. Thus we do not change the result of APPROX, and therefore REFINE returns. Since we do not (yet) have a partition where all cells are singletons (see Line 3 of the algorithm), we move on to calling SPLIT. The splitter uses the partition P_1 and branches on one of its cells. In theory, the ordering of the values returned by SPLIT does not matter, but it can make a significant difference in terms of practical performance.

We choose a shortest non-singleton cell in the partition, which at this stage means that we split at the cell [3,4,5]. This results in three graphs, each with one vertex coloured (individualising the corresponding point). Each of these graphs gives a child node, and a recursive call to SEARCH.

Down each branch, we go through REFINE again. Since G is 3-transitive, the refiners for these point stabilisers do not give graphs which change the result of APPROX. Continuing down the branch for the vertex 3, we call SPLIT again and colour 4 or 5, which always individualises both vertices (just ordered differently).

After this splitting, we have now fixed the list of vertices [8,3,4,5] (or [8,3,5,4] in the other branch). $G_{[8,3,4,5]}$ is the trivial group, and therefore the refiner R_G can now give every vertex a different colour. We represent this by ordering the vertices. However, the refiner must still be compatible, meaning this order must be compatible between different branches.

In our implementation, we choose an ordering for the orbits by using the minimal image of the list [8,3,5,4] of fixed points under the action of G . In this case, we find a permutation that maps [8,3,4,5] to its minimal image [1,2,3,6] under G . For example we can use the permutation (1,4,3,2,5,6,7,8). This maps the set of non-fixed points {1,2,6,7} to the set {4,5,7,8}. We then pick an ordering for the orbits in the minimal image, for example by length, and then break ties by smallest member of the orbit, and then map the list of orbits back. The ordering in the minimal image is therefore [4,5,7,8]. Applying the inverse of our minimising element above, we obtain a full ordering of the vertices and hence an ordered partition: [8|3|4|5|1|2|6|7]. When we return from REFINE, PART fixes every point, and hence we generate a minimal image under G of these fixed points, with minimising permutation (1,4,3,2,5,6,7,8). We used this permutation earlier – this will not always happen, but since we use minimal image in many places, the same permutations often appear. The present node is a leaf, and we return.

Going back to the node where we split for the second time, there is another branch where we colour first 3, then 5, then 4. There we reach a leaf in the same way, with the minimising permutation (1,7,4,5,3,2,8) mapping [8,3,5,4] to [1,2,3,5]. The corresponding ordered partition is [8|3|5|4|7|6|1|2]. Going back one step further, we see that the remaining two branches also split once more, and in total we have six leaves. For completeness, we describe the partitions that belong to the remaining leaves:

If we go down the branch where we first fix 4, then 3, then 5, then the minimal image of [8,4,3,5] under G is [1,2,3,5], with a minimising element (1,7,8)(2,6,4). The final partition into singleton cells, at the leaf, is [8|4|3|5|6|2|1|7]. Fixing first 4, then 5, then 3 gives minimal image [1,2,3,6] with minimising permutation (1,4,2,7,5,3,6,8). The partition at the leaf is [8|4|5|1|7|3|2|6]. For the branch where we fix first 5, then 3, then 4, we minimise [8,5,3,4] with the element (1,4,6,5,2,8), which gives the partition [8|5|3|4|1|6|7|2] at the leaf. Finally, if we fix 5, 4 and 3 in this order, then the minimising element is (1,7,6,8)(2,4,3,5) and the partition at the leaf is [8|5|4|2|3|7|1|6]. Therefore, in total, our search generates six minimising permutations and hence six candidates for a canonical image.

Note that the execution of Algorithm 1 is compatible with the action of G on ACTSET. This is because each sub-function we execute during the algorithm is, by construction, compatible with the action of $\text{Sym}(\Omega)$, except the refiner R_G . But this is compatible with the action of G .

Remark 3.3. Suppose that $g \in \text{Sym}(\Omega)$, that $a, b \in \text{ACTSET}$ and that $a^g = b$. Moreover let $n, k \in \mathbb{N}$ and let $[R_1, \dots, R_n]$ be a list of compatible refiners for ACTSET, $[\text{SPLIT}_1, \dots, \text{SPLIT}_k]$ a list of compatible splitters and APPROX a compatible approximator for producing a search tree for a as in Algorithm 1.

Then mapping T_a to T_a^g gives the same tree as producing T_b by following Algorithm 1.

We point out once more that, for this remark to hold, it is important that the children of nodes are treated as unordered in theory, although they will have an ordering in the implementation.

4. Finding canonical images using search trees

In this section we describe how to find a canonical image from a search tree, and we prove that what we find satisfies all conditions of a canonical image. Rather than providing a mathematical description of the canonical image of an object a , our proof will instead show directly that, if two objects a and b are in the same orbit under some action, then they will have the same canonical image. Much of this proof will come from *compatibility*, and analysing the relationship between the search trees for a and b .

Another important requirement is that our algorithm can be implemented efficiently. We will mention briefly where the algorithm is designed to allow efficient implementation, at the cost of complicating the description. Further, at the end of this section we will discuss some improvements which greatly improve performance, without changing the canonical image.

Remark 4.1. Suppose that $G \leq \text{Sym}(\Omega)$ and that l is a list that contains each element of Ω exactly once. Then there is a unique permutation, which we call a **minimising permutation**, in G that maps l to its minimal image under G with respect to the lexicographic ordering on lists. The minimal image of l and the minimising permutation can be calculated in polynomial time (see appendix).

Example 4.2. Let $\Omega := \{1, \dots, 8\}$ and $G := \langle (1,2,3,4,5,6,8), (1,3,2,6,4,5), (1,6)(2,3)(4,5)(7,8) \rangle \leq \text{Sym}(\Omega)$. We consider the minimal images of the lists [8,3,4,5,1,2,6,7] and [8,3,5,4,1,2,6,7] with respect to lexicographical ordering of lists using the natural ordering of Ω – under both $\text{Sym}(\Omega)$ and G . Since $\text{Sym}(\Omega)$ is 8-transitive on Ω , the minimal image of both lists under $\text{Sym}(\Omega)$ is [1,2,3,4,5,6,7,8], with minimising permutations (1,5,4,3,2,6,7,8) and (1,5,3,2,6,7,8) respectively. Under G we obtain the minimal images [1,2,3,6,4,5,7,8], with minimising permutation (1,4,3,2,5,6,7,8), and [1,2,3,5,7,8,6,4] with minimising permutation (1,7,4,5,3,2,8), respectively.

We will now explain the connection between our search trees and canonical image.

Definition 4.3. Let $G \leq \text{Sym}(\Omega)$, \leq denote a compatible partial order and let PART denote an ordered orbit approximator.

We define a function CANONICAL from ACTSET to itself by following the steps below. We will prove later that this actually is a function, and that it maps members of ACTSET to a canonical member of their orbit under G . Let $a \in \text{ACTSET}$ and let T_a be the search tree for a as defined in Algorithm 1. Recall that this means that each node is labelled with a *graph stack*, and leaf nodes are also labelled with a *result* permutation.

1. Denote the set of leaves of T_a as L_a .
2. For each leaf $l \in L_a$ define a list of graph stacks $\text{STACKS}(l)$ by taking, in order, the *graph stack* vertex label of each vertex on the path from the root of T_a to l . We remark that each member of this list will be a prefix of the next one, because graph stacks accumulate graphs as we go along a branch through the search tree.

3. Extend the partial order \leq to the elements of $\{\text{STACKS}(l) \mid l \in L_a\}$ by ordering these lists lexicographically, comparing pairs of elements by \leq .
4. Define $M_a := \{m \in L_a \mid \text{There is no } l \in L_a \text{ such that } \text{STACKS}(l) < \text{STACKS}(m)\}$. This means that M_a consists of those elements of L_a which have minimal images with respect to \leq when STACKS is applied.
5. Build the set of permutations $P_a = \{\text{result}(m) \mid m \in M_a\}$, where $\text{result}(m)$ is the result permutation label from the leaf node m .
6. Finally, define $\text{CANONICAL}(a) := \text{Min}_{\leq} \{a^g \mid g \in P_a\}$.

Note that the definition of STACKS in Definition 4.3 as a list of graph stacks is chosen in this way to allow efficient algorithm implementation. The algorithm would still produce a correct canonical image if we only compared the graph stacks from the leaf nodes in steps 3 and 4 of Definition 4.3, and this would also simplify describing and proving the correctness of the algorithm. The advantage of using STACKS is that it allows us to sometimes prove a node is not in M_a in Step 4 by only considering a prefix of STACKS .

Lemma 4.4. *Let CANONICAL be as defined in Definition 4.3. Then CANONICAL is a well-defined function from ACTSET to itself, and moreover, it is a canonical labelling function with respect to the action of G on ACTSET .*

Proof. Let $a \in \text{ACTSET}$ and suppose that T_a and T'_a are search trees constructed as in Algorithm 1. Both trees start with the same root node, and the same refiners are applied in the same ordering. A choice occurs only when a splitter is applied, because the children could appear in different orders for T_a and T'_a . This is relevant for implementation, but in theory the trees do not differ, so we obtain the same set L_a of leaves in both cases. This also means that the set $\{\text{NODES}(l) \mid l \in L_a\}$ is independent of whether we build the tree T_a or the tree T'_a . Since we started with the same ordering \leq and hence the same induced lexicographical ordering, it follows that the sets M_a and P_a are also independent of the choice of the search tree. This implies that $\text{CANONICAL}(a)$ is well-defined.

Next suppose that $b \in a^G$ and let $g \in G$ be such that $a^g = b$. Then Remark 3.3 tells us that T'_a is the same tree as if we produced T_b following Algorithm 1.

Hence the set of leaves L_b of T_b is the same as L_a^g , and for each $l \in L_b$ there is a leaf $t \in L_a$ such that $t^g = l$, so $\text{STACKS}(l) = \text{STACKS}(t)^g$. Since \leq is compatible with applying g , we also know that $M_b = M_a^g$. Moreover, PART satisfies $\text{PART}(a^g) = \text{PART}(a)^g$ and MINPERM satisfies the condition that $\text{MINPERM}(G, L^g) = \text{MINPERM}(G, L)^g$, and then it follows that $P_a \cdot g = P_b$. In our final step we consider the minimal members of P_a and P_b :

$$\begin{aligned} & \{a^{p^{-1}} \mid p \in P_a\} \\ = & \{a^{gsg^{-1}p^{-1}} \mid p \in P_a\} \\ = & \{a^{g(pg)^{-1}} \mid p \in P_a\} \\ = & \{a^{gq^{-1}} \mid qg^{-1} \in P_a\} \text{ (where } q = pg\text{)} \\ = & \{a^{gq^{-1}} \mid q \in P_a \cdot g\} \\ = & \{b^{q^{-1}} \mid q \in P_b\} \end{aligned}$$

Therefore we have mapped a and b to a common set of images, namely $\{a^{p^{-1}} \mid p \in P_a\} = \{b^{p^{-1}} \mid p \in P_b\}$. Given the definition of the canonical image of x as $\text{min}\{x^{p^{-1}} \mid p \in P_x\}$, for a fixed total ordering min on our objects, we deduce that $\text{CANONICAL}(a)$ and $\text{CANONICAL}(b)$ are equal. \square

Example 4.5. We continue the discussion from Example 3.2. There we had six leaves, and we consider Definition 4.3 for each of them. Starting with the leaf $l_1 := [8, 3, 4, 5, 1, 2, 6, 7]$, we follow the path from the root node to this leaf as follows:

We begin at the root node, with the empty graph stack. The first child has two graphs, created by the refiners, namely the graph Γ itself

from the refiner for Γ , and the complete graph C_Ω on Ω (from refining G by orbits, because G is transitive).

At this point we have completed the first node, and therefore, this will be the first member of STACKS for every leaf node. For the first child, we branch on Γ_3, Γ_4 and Γ_5 in turn, and in each branch we create a new node. Then we build a graph where vertices have the same colour if and only if they are in the same orbit of $G_{[8,3,4,5]}$, which means that every vertex is individualised. This last graph will be abbreviated with l_1 itself, the singleton partition corresponding to it, because the ordering of the individualisation will matter in practice. We conclude that $\text{STACKS}(l_1)$ looks like this:

$$[\emptyset, [\Gamma], [\Gamma, C_\Omega], [\Gamma, C_\Omega, \Gamma_3], [\Gamma, C_\Omega, \Gamma_3, \Gamma_4], [\Gamma, C_\Omega, \Gamma_3, \Gamma_4, \Gamma_5], [\Gamma, C_\Omega, \Gamma_3, \Gamma_4, \Gamma_5, l_1]].$$

The i^{th} member of STACKS is the graph that was added to the stack at depth i , if we count from the root node. Note that in practice this is not stored as a list of stacks, but as a single stack, and the size of the stack at the end of each node. Also, we assume that refiners do not add duplicate graphs – adding these would not change the results or the correctness of the algorithm (as long as the graphs were added consistently).

Using the same notation and looking at the remaining leaves in the ordering in which we describe them in Example 3.2, the remaining lists coincide with $\text{STACKS}(l_1)$ up to the 3rd entry. Thus, we do not repeat these entries and only give the remaining ones.

$$\begin{aligned} \text{STACKS}(l_2) &= [\dots, [\Gamma, C_\Omega, \Gamma_3], [\Gamma, C_\Omega, \Gamma_3, \Gamma_5], [\Gamma, C_\Omega, \Gamma_3, \Gamma_5, \Gamma_4], [\Gamma, C_\Omega, \Gamma_3, \Gamma_5, \Gamma_4, l_2]], \\ \text{STACKS}(l_3) &= [\dots, [\Gamma, C_\Omega, \Gamma_4], [\Gamma, C_\Omega, \Gamma_4, \Gamma_3], [\Gamma, C_\Omega, \Gamma_4, \Gamma_3, \Gamma_5], [\Gamma, C_\Omega, \Gamma_4, \Gamma_3, \Gamma_5, l_3]], \\ \text{STACKS}(l_4) &= [\dots, [\Gamma, C_\Omega, \Gamma_4], [\Gamma, C_\Omega, \Gamma_4, \Gamma_5], [\Gamma, C_\Omega, \Gamma_4, \Gamma_5, \Gamma_3], [\Gamma, C_\Omega, \Gamma_4, \Gamma_5, \Gamma_3, l_4]], \\ \text{STACKS}(l_5) &= [\dots, [\Gamma, C_\Omega, \Gamma_5], [\Gamma, C_\Omega, \Gamma_5, \Gamma_3], [\Gamma, C_\Omega, \Gamma_5, \Gamma_3, \Gamma_4], [\Gamma, C_\Omega, \Gamma_5, \Gamma_3, \Gamma_4, l_5]], \text{ and} \\ \text{STACKS}(l_6) &= [\dots, [\Gamma, C_\Omega, \Gamma_5], [\Gamma, C_\Omega, \Gamma_5, \Gamma_4], [\Gamma, C_\Omega, \Gamma_5, \Gamma_4, \Gamma_3], [\Gamma, C_\Omega, \Gamma_5, \Gamma_4, \Gamma_3, l_6]]. \end{aligned}$$

The final entries l_2, \dots, l_6 correspond to the remaining singleton partitions from Example 3.2. In our ordering of graphs, the graphs $\Gamma_3, \dots, \Gamma_8$ are considered to be the same because they have in common that only a single vertex is coloured and there are no arcs or edges.

We now attempt to order the different STACKS – this is done using a permutation-invariant ordering, in lexicographic ordering. Looking at $\text{STACKS}(l_2)$ and $\text{STACKS}(l_3)$, the identity permutation maps the first three entries (which we omitted) of $\text{STACKS}(l_2)$ to the corresponding entries of $\text{STACKS}(l_3)$, the permutation $(3, 4)$ maps the sixth entry to the sixth and $(3, 4, 5)$ maps the sixth and seventh entries simultaneously. This means that the search could not stop before reaching depth 8, because we cannot distinguish between the lists of graphs built at depths 1 to 7.

This means that in terms of the lexicographical ordering, the six stacks above only differ at the very last graph stack, and there only at the last entry. These final graphs are written as lists, which means that they can be ordered lexicographically. We see that l_1 is the smallest one, therefore $M_a = \{l_1\}$ and $P_a = \{(1, 4, 3, 2, 5, 6, 7, 8)\}$. Then the canonical image of the graph Γ is defined to be its image under $(1, 4, 3, 2, 5, 6, 7, 8)$, which is the graph with vertex set Ω and the following edges: $\{4, 7\}, \{4, 8\}, \{5, 7\}, \{5, 8\}, \{2, 1\}, \{3, 1\}, \{6, 1\}$.

4.1. Improving performance

When running our algorithm in practice, we improve performance by skipping the generation of parts of the search tree when we can prove those parts of the search tree will not contain the canonical image. The techniques we use are similar to those used in other backtrack search algorithms, in particular Nauty and our previous work on Graph Backtracking.

The most significant technique involves skipping generating the children of nodes if we can prove that the leaves generated from this node are not elements of the set M_a from Step 4 of Definition 4.3. Those parts of the tree which do not generate elements of M_a do not need

to be generated at all, during search, because they are not used. In practice, search trees are generated as a depth-first transversal, and we keep track of the current smallest list of graph stacks generated by any leaf. If some node of the tree produces a lexicographically larger list of graph stacks than the smallest known list of graph stacks for a leaf, then all children of this node will also create lexicographically larger lists of graph stacks. Therefore, generation of the entire sub-tree under this node can be skipped, because the ordering of lists of graph stacks in Definition 4.3 uses lexicographic ordering.

5. Experiments

In this section, we will measure the performance of our algorithm compared to alternatives. As there is no other implementation of a general framework for finding canonical images in arbitrary permutation groups, each experiment will, where appropriate, compare against specialised algorithms. All of our experiments are performed with GAP 4.12 [6]. Our algorithm is implemented in the Vole package [5]. For all instances, there is a time-out of 30 minutes and a memory limit of 4 GB. We note that we specifically search for difficult problems, both for our algorithm and existing algorithms.

5.1. Grid group experiments

As in our previous work on minimal and canonical images [8], we will use “grid groups”. We use these because, in practice, they often produce difficult problems. We reproduce the definition of a grid group here.

Definition 5.1. Let $n \in \mathbb{N}$. The direct product $\text{Sym}(n) \times \text{Sym}(n)$ acts on the set $\{1, \dots, n\} \times \{1, \dots, n\}$ of pairs in the following way:

For all $(i, j) \in \{1, \dots, n\} \times \{1, \dots, n\}$ and all $(\sigma, \tau) \in \text{Sym}(n) \times \text{Sym}(n)$ we define

$$(i, j)^{(\sigma, \tau)} := (i^\sigma, j^\tau).$$

The subgroup $G \leq \text{Sym}(\{1, \dots, n\} \times \{1, \dots, n\})$ defined by this action is called the $n \times n$ **grid group**.

We begin by reproducing the experiments of our minimal and canonical images paper [8], looking for the canonical image of randomly generated sets of points in an $n \times n$ grid group. Here, we consider sizes of grids from 5 to 50 inclusive (46 sizes of grids in total) and we note that a stabilizer chain for the group is calculated before starting the experiment. Otherwise the time is dominated by the time taking to build a stabilizer chain.

In [8] we compared 16 search orderings. We consider the very best one (RareOrbitPlusRare), and MinOrbit, as this most closely matches the search ordering Vole uses. Overall we considered 138 experiments, meaning that Vole solved every instance we considered within the allowed time limit. The one experiment where RareOrbitPlusRare failed reached the memory limit of 4 GB. The table below gives the number of tests which were solved in the time and memory limit, and the total time (in seconds) that was taken by those experiments which finished. We note that this means that algorithms which only finished few tests may also have a low time.

Size of set	$\lfloor \frac{n^2}{2} \rfloor$		$\lfloor \frac{n^2}{4} \rfloor$		$\lfloor \frac{n^2}{8} \rfloor$		total	
Results	solved	time	solved	time	solved	time	solved	time
Vole	46	53	46	58	46	106	138	218
RareOrbitPlusRare	45	185	46	18	46	10	137	214
MinOrbit	8	60	13	260	21	116	42	438

Vole is overall the same speed as the RareOrbitPlusRare algorithm. This is despite the fact that it contains no specialised code for finding canonical images of sets of points and can also be used to find canonical

images of more complex combinatorial structures, such as sets of sets, or sets of tuples, which are not supported by the algorithm in [8].

Next we looked at canonical images of permutations under conjugation in the grid group. Since there is no other tool available for this purpose, as far as we know, we compared the time it took to find the canonical image of two permutations, g and h , under a group G against the time it takes GAP to find a permutation which can map g to h in G , known as `RepresentativeAction`. Since the canonical images of g and h can be used to implement `RepresentativeAction`, this is a reasonable comparison.

However, this can be viewed as slightly unfair, with an advantage for Vole. If we wanted to sort n permutations into conjugacy classes, for example, then we would only have to call Vole n times, whereas GAP may require calling `RepresentationAction` on all pairs of permutations.

We chose to only use permutations built from pair-wise disjoint 2-cycles (i.e. involutions). When the size of cycles grows, and when there are cycles of many different sizes, then both Vole and GAP solve problems very quickly and the time taken becomes dominated by the time taken to build an initial stabiliser chain for the group. For each size of grid, we consider ten randomly generated permutations. The results show both the number of problems solved and the average time taken for the solved instances in seconds.

n	GAP		Vole	
	# solved	time (sec)	# solved	time (sec)
10	10	0.9	10	0.7
20	1	58	10	13
30	0	-	10	101
40	0	-	10	448

We can see that Vole is much faster than GAP and able to solve problems far beyond the ability of GAP.

Out of curiosity, we also looked at the performance of MAGMA (see [3]) in checking involutions in grid groups for conjugacy. MAGMA solved this problem much faster than either GAP, or Vole. It solved the problems in 50×50 grid groups in 0.5 to 7 seconds. According to the documentation for MAGMA, this functionality is implemented using Leon’s Partition Backtrack [14,15], while greatly outperforming the implementations in GAP and Vole – at least when we only need to check if two permutations are conjugate. This is a constant reminder for us that further improvements are possible and necessary. We were unable to find out why MAGMA is able to solve these problems so quickly, compared to the implementation of Leon’s Partition backtrack code in GAP. For example, it would be interesting to know whether the impressive speed is possible because of the MAGMA function `CLASS-REPRESENTATIVE`. If that was used and was at least partly responsible for the performance, then our work on canonical images might pave the way for future speed-ups in GAP, and in other systems as well.

5.2. Canonical images under the action of primitive groups

In this experiment we reproduce an experiment from [8], where we look for canonical images of randomly generated sets under the action of a primitive group which moves between 10 and 300 points. We remove the natural alternating and symmetric groups, because these groups can easily be treated as special cases.

The results are given below, showing both the best results for static and dynamic selection strategies from [8], and Vole. The total number of instances solved, and the time in seconds taken for these solvable problems is given. A total of 5,952 experiments were attempted, meaning no system completed every problem.

While Vole does take longer on average, it solves a larger number of instances. The most significant reason for the slowdown is that Vole is generating orbital graphs for the groups as search progresses. In some

Size of sets	$\left\lfloor \frac{n}{2} \right\rfloor$		$\left\lfloor \frac{n}{4} \right\rfloor$		$\left\lfloor \frac{n}{8} \right\rfloor$		total	
	# solved	time (sec)	# solved	time (sec)	# solved	time (sec)	# solved	time (sec)
Vole	1963	7517	1984	8573	1984	2410	5931	18501
RareOrbitPlusRare	1956	2653	1983	627	1984	46	5923	3327
FixedMinOrbit	1916	4139	1979	907	1984	69	5879	5116

cases these graphs are very large and do not reduce the size of the search. In the future we plan to look for better heuristics to filter which graphs are generated and used during refinement. Also, in this experiment, we did not re-use the graphs on different calls to Vole.

5.3. Canonical images of graphs

Vole can solve a much wider range of problems than previous graph canonising systems such as Nauty/Traces [17] and Bliss [13]. This is because it supports finding canonical images under any group of symmetries on the vertices of the graph, while Nauty, Traces and Bliss can only search within the full symmetric group on the vertex set.

Here, we therefore compare Vole and Nauty on the canonical image problems that Nauty can solve. We use the latest version of Nauty, and we found that this was faster than Bliss for these experiments.

On these problems, there is no fundamental reason that Vole should not perform similarly to Nauty. However, in practice, it is significantly slower in terms of the fastest speed Vole can be, a slowdown of around 5 times. After investigating how our algorithm performed, we believe that there are three main reasons for this:

- Vole supports refining stacks of graphs, and the code is less efficient when only a single graph is refined.
- Vole uses stabiliser chains to build the known stabiliser as the search progresses – this calculation is done in GAP and therefore requires communication costs.
- Nauty uses a different, better ordering to decide which cell to split on.

In general, Nauty has been carefully optimised for over 40 years, while Vole is comparatively new, and hence there are many small optimisations which still have to be considered and implemented. We believe there is no optimisation in Nauty which could not, in principle, be added to Vole.

Vole can solve graph problems which Nauty cannot – finding canonical images under the action of groups other than the symmetric group on the set of vertices. For example, we look for the canonical image of graphs under the action of cyclic groups and alternating groups on the whole set of vertices. In our experiment, we compare Nauty and Vole in the symmetric group, because Nauty can only compute canonical images of graphs under the full symmetric group. We also use Vole to search for canonical images of graphs in the cyclic group and alternating group on the set of vertices – showing how Vole can solve interesting problems that have been impossible to solve with previous graph automorphism systems such as Nauty.

The slowdown that we observe when we search for canonical images under subgroups, rather than under the full symmetric group, is entirely due to calculating stabiliser chains. The only reason why it does not occur for the symmetric groups is that they are treated as a special case. In our experiments, the searches in smaller groups explore the same number of nodes, or even fewer nodes. This behaviour is not guaranteed, but it usually occurs in practice. We could add similar optimisations for the cyclic and alternating groups, but in general our algorithm’s performance is greatly affected by the performance of stabiliser chain calculations, which is why we plan to improve these calculations in the future. In particular, currently GAP, and therefore Vole which re-uses GAP stabiliser chain implementation, cache at most a single stabiliser chain at a time, producing repeated work if a previously created stabiliser chain is needed again.

For our experiments we use the “Steiner Triple System Graphs” from the website of Nauty and Traces, which were originally distributed with Bliss (<https://pallini.di.uniroma1.it/Graphs.html>). The times in these experiments are given in seconds.

instance	Nauty (symmetric)	Vole (symmetric)	Vole (cyclic)	Vole (alternating)
sts-13	0	0.014	0.055	0.099
sts-15	0	0.016	0.071	0.309
sts-19	0.001	0.031	0.134	1.235
sts-21	0	0.034	0.184	1.994
sts-25	0.009	0.086	0.329	9.398
sts-27	0.001	0.099	0.423	18.848
sts-31	0.035	0.251	0.869	51.666
sts-33	0.001	0.233	1.019	133.580
sts-37	0.089	0.601	1.441	265.022
sts-39	0.004	0.632	1.953	623.571
sts-43	0.219	1.409	2.918	1124.527

5.4. Results overview

Overall, we can see that Vole has the potential to become competitive with previous work on finding canonical images, after some optimization, with the exception of Nauty. Also, Vole can solve a wide range of other canonical image problems which were previously insolvable by any system, and these can be solved very efficiently.

6. Conclusion

In this paper we present a general framework for finding canonical images, and demonstrate its practical value on a range of experiments. Compared to previous systems, our implementation Vole can solve a wider range of different canonical image problems.

One significant advantage of our framework is that the same refiners can often be used to solve intersection and coset problems. We hope this will create a virtuous circle, as new refiners will have, in many cases, immediate applications to both intersection and canonical image problems. Almost all refiners in this paper were re-used, unchanged, from our earlier work – only one new refiner (described in 7.1) was created, and that refiner is now the refiner Vole uses for a group given by a list of generators. In the future we plan to work on refiners for normalisers, which based on the framework in this paper will also be used for finding canonical images of groups under conjugacy.

7. Additional results and comments

Here we present two algorithms which are needed in this paper. Firstly, we present how to find the minimal image of a list of elements of Ω with respect to the action of a group $G \leq \text{Sym}(\Omega)$. Secondly, we present a new refiner for a group given as a list of generators. Previously presented refiners for a group given as a list of generators could return the cells of partitions in different orders in different searches. This was not an issue previously, but it is for the algorithm we present in this paper.

In this section we say that a list of elements of Ω is **exhaustive and non-repetitive** if it contains each element of Ω exactly once.

Lemma 7.1. *Given a set Ω with a total ordering \leq , and a group $G \leq \text{Sym}(\Omega)$, let MINPERM_G denote the function that maps each non-repetitive exhaustive list on Ω to its unique minimising permutation in G .*

Then the images under MINPERM_G can be calculated in polynomial time.

Proof. If we have two exhaustive, non-repetitive lists on Ω , then there is a unique permutation in $\text{Sym}(\Omega)$ that maps one to the other. Therefore MINPERM_G is a well-defined function.

Given a list L with members drawn from Ω , the minimal image of L under G can be found in a series of steps. Firstly, the minimal image under a lexicographic ordering must map the first member $L[1]$ of L to its smallest value m_1 under G with respect to the ordering \leq on Ω . Finding a permutation $g \in G$ which maps $L[1]$ to m can be done by finding the orbit of $L[1]$ under G . The minimal image of L under G is the same as the minimal image of L^g under the point stabilizer G_{m_1} , because in the minimal image of L^g , m_1 must be fixed. We can continue mapping each element of L to its minimal image under the subgroup of G which maps all previous members of L to their smallest values, respectively. This calculation is finished after finitely many steps. Calculating orbits and point stabilizers in permutation groups can be performed in polynomial time using stabiliser chains, and consequently the whole calculation can be done in polynomial time. \square

7.1. Refiners for groups defined by minimal images

Searching for canonical images inside a group G requires a refiner for G . Typically G will be specified as a subgroup of a larger permutation group, with a set of generators. A family of refiners for a group given by a set of generators is given in Section 7.3.2 of [10], and it uses the orbital graphs of the group. This extends a previous refiner given in [7]. Leon also gives a refiner which uses only the orbits of the group [15]. These algorithms are different to any other refiners presented in the literature, as they do not describe a single refiner – instead they describe how to build a refiner dynamically during search. This means that the refiners can produce different outputs when solving different problems.

These refiners are not suitable for use in our canonical image algorithm, because the proof of correctness requires that a single fixed refiner is used for G . In this section we present a new refiner for groups given by a list of generators. The output of this refiner is almost identical to the previously presented refiner, the only difference is that the output is fixed, and therefore it produces the same output in every search, given the same input graph stack.

We begin by discussing the most important pieces of our refiner, and then we use this to describe the problem previous refiners were trying to solve and to explain our new solution for this problem.

Definition 7.2. Given a group G on a set Ω , a **fixed point refiner** F for G is a function from the set of lists of points of Ω to $\text{DIGRAPHSTACKS}(\Omega)$ such that for each list L of points from Ω and $g \in G$ the following holds: If $L^g = L$, then $F(L)^g = F(L)$.

Note that a fixed point refiner, as given in Definition 7.2, does not define a valid refiner in the sense of Definition 2.3 and [9] for two reasons – its input is a list of fixed points, rather than a graph stack, and it defines only a single function, rather than a pair of functions.

Example 7.3. Given a group G on a totally ordered set Ω , we give two examples for fixed point refiners. In both cases we make use of G_L , the point-wise stabilizer in G of a list L of points of Ω .

The Orbit refiner **FORBIT** maps a list L to a graph on Ω with no edges, where the vertices are labelled with the smallest member of Ω in their orbit under G_L .

The Orbital refiner **FORBITAL** maps a list L to a graph stack, where each graph is an orbital graph of G_L . These graphs are ordered in lexicographical order, using the total ordering of Ω .

The orbit refiner from Example 7.3 is similar to the technique used by Leon [15], while the orbital refiner is the technique used in [7] and [10]. The most obvious way to turn a fixed point refiner into a refiner in the sense of [9] would be to take the fixed points of the graph stack using a fixed point approximator, and then apply the fixed point refiner to its output. However, this obvious way does not work: Given a refiner

(R, R) for a group G (refiners for groups always consist of a pair of identical functions), it must satisfy the condition that for all $g \in G$ and all graph stacks S , it is true that $R(S^g) = R(S)^g$. The way we order the orbits and orbital graphs in Example 7.3 will not satisfy this, because the orbits and orbital graphs are sorted. In previous work this problem was solved by storing all the orbits, or orbital graphs, ever created. Then, if a list of fixed points L' was created which could be mapped to a previously created list L by some $g \in G$, then the image of the refiner output for L under g^{-1} could be returned.

The reason for the dynamic nature of this refiner is that it takes the first list of fixed points from each orbit as a “canonical member” of that orbit. We deal with this problem by explicitly choosing the “canonical member” we will use for each orbit in advance – the minimal image of the list of fixed points.

Lemma 7.4. Given a fixed point refiner F for a subgroup G of $\text{Sym}(\Omega)$ and a fixed point approximator **FIXED**, we define the function **REF** from $\text{DIGRAPHSTACKS}(\Omega)$ to itself as follows:

If $S \in \text{DIGRAPHSTACKS}(\Omega)$, $L := \text{FIXED}(S)$ and $g \in G$ is a permutation which maps L to its minimal image under G , then $\text{REF}(S) := F(L^g)^{g^{-1}}$.

Then (REF, REF) is a refiner for G in the sense of [9].

Proof. Firstly, we prove that **REF** is well-defined, as the definition defines g as any permutation in G which maps L to its minimal image under G . For this let L be a list of points in Ω and let $g, h \in G$ be such that $M := L^g = L^h$ is the minimal image of L under G . Then $M^{g^{-1}h} = M$, so by the definition of a fixed point approximator, $F(M)^{g^{-1}h} = F(M)$. This means that $F(M)^{g^{-1}} = F(M)^{h^{-1}}$ and therefore $F(L^g)^{g^{-1}} = F(L^h)^{h^{-1}}$.

Next we show that **REF** is a valid refiner for G as explained in Definition 4.1 of [9]. For all $S, T \in \text{DIGRAPHSTACKS}(\Omega)$ such that $S \cong T$, we must prove that

$$G \cap \text{Iso}(S, T) \subseteq G \cap \text{Iso}(\text{REF}(S), \text{REF}(T)).$$

It suffices to show that for all $g \in G$, where $g \in \text{Iso}(S, T)$, g is also in $\text{Iso}(\text{REF}(S), \text{REF}(T))$.

Let $L_S := \text{FIXED}(S)$ and $L_T := \text{FIXED}(T)$. By the definition of a fixed point approximator, the fact that $S^g = T$ implies that $L_S^g = L_T$. Therefore L_S and L_T are in the same G -orbit, and therefore they have the same minimal image L_M . Further, let $h \in G$ be such that $L_T^h = L_M$ and note that this means that gh maps L_S to L_M .

Then $\text{REF}(S) = F(L_S^{gh})^{(gh)^{-1}} = F(L_M)^{(gh)^{-1}} = (F(L_M)^{h^{-1}})^{g^{-1}}$ and $\text{REF}(T) = F(L_T^h)^{h^{-1}} = F(L_M)^{h^{-1}}$. This shows that $g \in \text{Iso}(\text{REF}(S), \text{REF}(T))$, completing the proof. \square

In practice, implementing the refiner in Lemma 7.4 is very similar to implementing the refiners for a group given by a list of generators as previously described in [7,10,15]. They each involve finding permutations which map a list of fixed points to a “standard list” – the difference is that, in the previous algorithms, this standard list was chosen as the first element of each orbit to occur, while now it is fixed as the minimal image of each orbit. The old and new refiners will, when using the same general technique (either orbits, or orbital graphs), produce the same graph stacks, up to re-ordering of the graphs and re-naming the labels used for vertices and edges. We benchmarked this new algorithm against implementations of the previous refiners and found no measurable difference.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Rebecca Waldecker reports financial support was provided by Volkswagen Foundation. Chris Jefferson reports financial support was provided by The Royal Society. Rebecca Waldecker reports financial support was provided by German Research Foundation.

Data availability

We have uploaded files for the experiments.

Acknowledgements

Many of the ideas that lead to this work stem from our project on Graph Backtracking, which is why we thank the Volkswagen-Stiftung (Grant no. 93764) and the Royal Society (Grant code URF\R\180015) for their financial support of various aspects of this project. The more recent work has been supported by the DFG (Grant no. WA 3089/9-1) and again the Royal Society (Grant codes RGF\EA\181005 and URF\R\180015). We are grateful to all these institutions for their support. Also, we would like to thank the referees for many valuable comments and suggestions on an earlier version of this article.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jaca.2023.100010>.

References

- [1] László Babai, Graph isomorphism in quasipolynomial time, <https://arxiv.org/abs/1512.03547>, 2018.
- [2] Anton Betten, The orbiter ecosystem for combinatorial data, in: ISSAC'20 – Proceedings of the 45th International Symposium on Symbolic and Algebraic Computation, ACM, New York, 2020, pp. 30–37 (English summary).
- [3] Wieb Bosma, John Cannon, Catherine Playoust, The Magma algebra system. I. The user language, *J. Symb. Comput.* 24 (1997) 235–265.
- [4] Peter Cameron, *Permutation Groups*, Cambridge University Press, 1999.
- [5] Mun See Chang, Christopher Jefferson, Wilf A. Wilson, Vole – GAP package, Version 0.5.3, <https://peal.github.io/vole>, 2021.
- [6] The GAP Group, *GAP – Groups, Algorithms, and Programming*, Version 4.12.0, 2022.
- [7] Christopher Jefferson, Markus Pfeiffer, Rebecca Waldecker, New refiners for permutation group search, *J. Symb. Comput.* 92 (2019) 70–92.
- [8] Christopher Jefferson, Eliza Jonauskite, Markus Pfeiffer, Rebecca Waldecker, Minimal and canonical images, *J. Algebra* 521 (2019) 481–506, <https://doi.org/10.1016/j.jalgebra.2018.11.009>.
- [9] Christopher Jefferson, Markus Pfeiffer, Rebecca Waldecker, Wilf A. Wilson, Permutation group algorithms based on directed graphs, *J. Algebra* 585 (2021) 723–758, <https://doi.org/10.1016/j.jalgebra.2021.06.015>.
- [10] Christopher Jefferson, Markus Pfeiffer, Rebecca Waldecker, Wilf A. Wilson, Permutation group algorithms based on directed graphs (extended version), arXiv e-prints, <https://arxiv.org/abs/1911.04783>.
- [11] Christopher Jefferson, Rebecca Waldecker, Wilf A. Wilson, Perfect refiners for permutation group backtracking algorithms, *J. Symb. Comput.* 114 (2023) 18–36, <https://doi.org/10.1016/j.jsc.2022.04.007>.
- [12] Christopher Jefferson, Wilf A. Wilson, GraphBacktracking – GAP package, Version 0.5.2, <https://peal.github.io/GraphBacktracking>, 2021.
- [13] Tommi Junttila, Petteri Kaski, Engineering an efficient canonical labeling tool for large and sparse graphs, in: Proceedings of the Ninth Workshop on Algorithm Engineering and Experiments and the Fourth Workshop on Analytic Algorithms and Combinatorics, 2007, pp. 135–149.
- [14] Jeffrey S. Leon, Permutation group algorithms based on partitions. I. Theory and algorithms, *J. Symb. Comput.* 12 (1991) 533–583.
- [15] Jeffrey S. Leon, Partitions, refinements, and permutation group computation, in: *Groups and Computation, II*, New Brunswick, NJ, 1995, in: DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 28, 1997, pp. 123–158.
- [16] Steve Linton, Finding the smallest image of a set, in: Proceedings of ISSAC 04, 2004, pp. 229–234.
- [17] Brendan D. McKay, Adolfo Piperno, Practical graph isomorphism II, *J. Symb. Comput.* 60 (2014) 94–112.
- [18] Pascal Schweitzer, Daniel Wiebking, A unifying method for the design of algorithms canonizing combinatorial objects, in: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, ACM, 2019, pp. 1247–1258.